

A Case Study: Developing a Remote, Rapid, and Automated Usability Testing Methodology for On-line Books

Jean Scholtz
National Institute of Standards and Technology
Bldg. 225, A 216
Gaithersburg, MD 20899
Jean.Scholtz@nist.gov

Abstract

The author describes a case study of a usability test of an on-line engineering statistical handbook. The case study was conducted with two goals: to determine usability issues affecting the use of the handbook and to determine if a kiosk-based testing methodology would be feasible to use in this type of test. The paper outlines the design of the experiment and describes the lessons learned, with respect to the design of on-line books and the design of kiosk-based tests for on-line books. The kiosk-based test was conducted twice. In the first instance, usability engineers were present to monitor the experiment. In the second condition, the design team monitored the experiment. The results from the two groups are compared to see if the condition produced an effect.

1. Introduction

Our mission at the National Institute of Standards and Technology (NIST) is to provide tools, techniques, and metrics to allow others to evaluate their work. A recent focus of our group in the Information Access and User Interfaces Division of the Information Technology Laboratory is on usability of web sites and web applications.

Usability work for web sites is problematic in many instances. The time frame in which sites are developed often does not allow for formal testing. A survey on web design conducted by Vora [7] found that while 70% of the respondents evaluated their web pages for usability either "all the time" or "almost all the time", a majority of the evaluations were informal. That is, colleagues and team members were the participants in the evaluation, rather than users more representative of the projected user population. Vora suggests that usability professionals may have been over represented in the survey as almost 50% of the respondents said they did some sort of formal usability testing.

Our research has been directed towards the development of tools and methodologies that are rapid, can be used remotely, and are as automated as possible [4,5] Remote usability testing has been a topic that usability professionals have been considering for several years [1]. Tools and techniques with these properties are particularly well suited to web development. The development time for web applications is short. Forty-three percent of respondents in Vora's survey said that the average time to develop a set of web pages was from a few weeks to a month. Almost 27% said that the average development time was only a few days. Just under 9% said that the average time was over 3 months. In order to fit into short development cycles usability testing must be rapid. Automation certainly facilitates rapid development.]. We have developed the WebMetrics suite of software tools [2] to allow developers of web sites to easily conduct usability evaluations both locally and remotely. We are working on producing automatic analysis for some of our tools and automating the construction of different types of usability testing.

Currently, there are three software tools available in our WebMetrics suite: WebSAT, WebCAT, and WebVIP. WebSAT does an automatic analysis of a web page and issues warnings about possible usability problems. Not all usability problems can be detected using WebSAT nor is every potential problem detected an actual problem. WebSAT documentation explains the checks that are done. This tool should be used as a starting point to determine the need for more usability testing.

WebCAT automates category analysis for web developers. The user can specify the categories used in a web site and the items located under each category. The test will be automatically created using this information. The test can then be distributed via the web and potential users can remotely be test participants. The data is automatically analyzed for the web developer.

WebVIP has both the remote and automated properties. Usability professionals can use this tool to "wrap" web

sites. This procedure produces a log of the links that users follow along with a time stamp of when the link is clicked. Usability tests can be run remotely by wrapping a site, contacting usability participants by e-mail, and providing them with a set of tasks to be carried out using the "wrapped" site. The log for each participant will be automatically generated and can be analyzed by the usability professional. We are currently working on an automated analysis component for this tool.

We have developed these tools by conducting case studies to determine the functionality and benefits of such tools. After the tool is developed, we make it publicly available. We hope to collect data from users of our tools, along with conducting more case studies ourselves, to determine the types of web sites and applications and time frames for which each tool is most useful. In addition, we are using these tools in our own work and using our experience to refine them. In the case study presented here, we used WebVIP to collect user path data.

2. The Case Study

We were asked if we would be willing to help in assessing the usability of an on-line (web based) statistical engineering handbook. This was extremely interesting to us as our previous research had focused on web sites and web applications. We were interested in usability issues for on-line books well as developing a methodology for usability testing methodologies for on-line books. We were involved early in the development of the book and so our work could affect navigation and structure of the book.

2.1. Why Choose a Kiosk-Based Evaluation?

The idea of doing a kiosk-based test originated during a discussion with the design group about obtaining participants for the usability test. The design team mentioned that they would be giving a demonstration of the on-line handbook during a conference that would be held at NIST. We decided to develop a short test that could be run during that conference as a companion to the demonstration. Since the intended users of this site were geographically distributed, we also wanted the capability of running additional tests remotely.

3. Kiosk-Based Testing

In this paper, we describe a case study and the implications for designing a tool to facilitate kiosk-based testing. The case study will be discussed both in terms of the lessons we learned about kiosk-based testing as well as the information we learned about design of on-line books. Although we specifically designed this

experiment for use as a kiosk, the study could also have been run remotely using the web.

3.1. The NIST/Sematech Handbook

The table of contents of the web-based handbook (at the time of the usability test) is shown in Figure 1. The handbook is a joint project between statistics researchers at NIST and Sematech. This handbook is integrated with statistical software so engineers can download the software and run the case studies contained in the handbook. They can also use the software to compute statistics for their own data.

The designers anticipate a wide variation of statistical knowledge in their audience; ranging from beginning engineers to experienced professionals. The handbook also needs to accommodate a range of browser skills. The designers wanted to provide navigation mechanisms within the document as well as the traditional browser navigation tools.

3.2. Kiosk Design

To date, there is little in the literature on designing kiosks. At CHI '89 conference attendees could use an InfoBooth kiosk sponsored by Apple Computer, and this author did. Salomon (1990) reflected on the design and use of this kiosk. Lessons learned included the following:

- Navigation aids within a given component should be easily distinguished from general navigation aids.
- Wording on controls and in directions is very important and should be pilot tested carefully.
- Information on how to use the kiosk should be clear and concise.
- The kiosk should use simple, easy to understand methods for navigating and should not rely on prior experience of the subjects.

3.3. What is Kiosk-based Testing?

Kiosk-based testing is a stand-alone test. The test can be placed in a kiosk, available in a public location, for anyone to use. Our ultimate public location choice would be the web. Because no human help is available for kiosk-based applications (e.g., ATMs) the interface and the tasks have to be completely self-explanatory. As we have no contact with users and no idea who the users are, there has to be a way to gather demographic data. Users of such a test do not want to spend a long time taking the test. The test has to be designed to be short, but yet collect meaningful data. The test instructions

have to be clear but also have to be well integrated into the material to be tested so that users do not have to spend time moving between directions and tasks. Users should also be allowed to quit at any time or to skip any portion of the test without affecting the rest of the test.

We based our design for the kiosk on Salomon's work plus the additional constraints that were added because the main goal of the kiosk was evaluation. To summarize, the guidelines for a kiosk-based usability test are:

- It can be completed in a short amount of time.

- The test is self-explanatory.
- The test directions are well integrated with the test materials.
- The test is designed so that any portion of it can be skipped and the test participant can continue with the remaining portions.

As in most designs, simplicity should be the guiding principle of kiosk-based evaluations. Simple wording, simple instructions, and limited sets of choices in the kiosk shell allow the user to concentrate on the evaluation tasks.

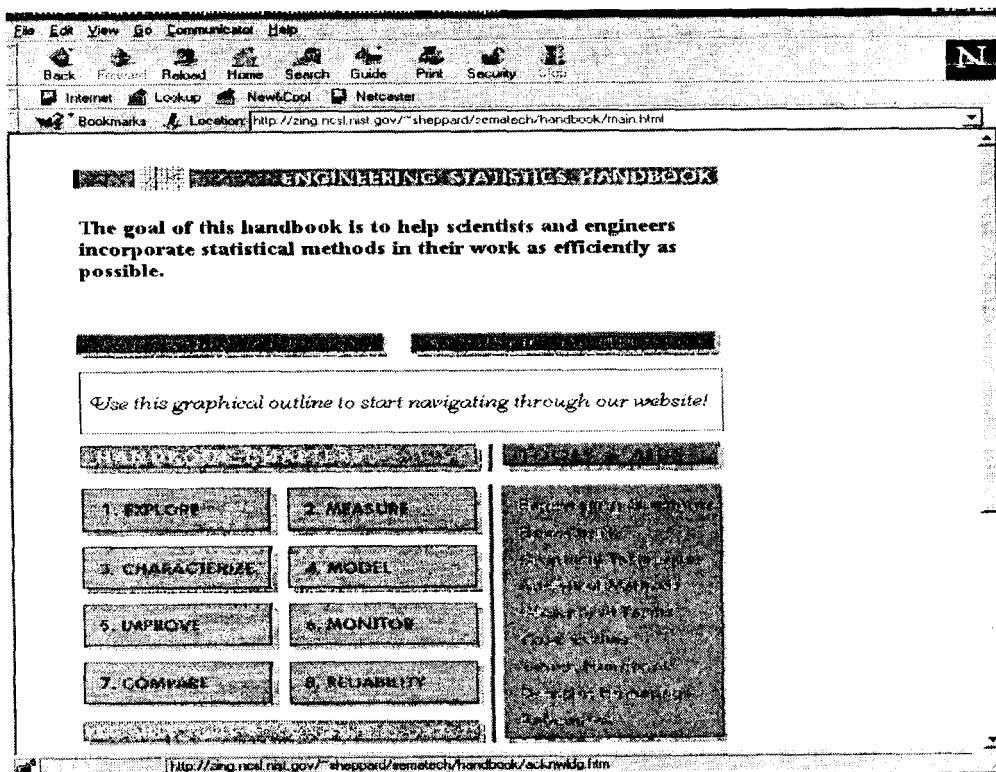


Figure 1. The first page of the NIST/Sematech Handbook

3.4. What Type of Applications Can be Tested in a Kiosk-based Test?

From our previous case studies, we had a good idea that regular web sites would be good candidates for this type of usability testing, as would small, focused applications. Testing a large document, like a handbook, would present challenges. One of our main concerns was how to conduct the test without requiring the participants to do an excessive amount of reading which would require much more time than subjects in a kiosk-based test would tolerate.

3.5. What type of usability information can be obtained in a kiosk-based test?

A self-contained test can provide us information about how a task is performed, the success of the task, and the time required to carryout the task. We can also collect feedback from participants. What we cannot capture is why participants performed the task in a given fashion and what terminology or affordances in the design were particularly helpful or confusing. Would the quantitative information be sufficient for our purposes? A kiosk-based test is an opportunity to collect data from many

users. Therefore, collecting quantitative information from a number of users would give an overall assessment of the usability of the handbook. If particular problems are discovered, smaller but more in-depth usability tests can be conducted to determine the particular aspects that are confusing and to arrive at alternative designs.

4. Designing the Experiment

In this section, we'll discuss the design of the experiment first from the perspective of the goals of the design team and secondly, from the perspective of the goals for our research in remote, rapid, and automated evaluation methods.

4.1. Goals of the Design Team

The design team had proposed a simple design for the table of contents of the book (see Figure 1). The chapter headings had been kept to one word and, while this made an elegant design possible, they were concerned about the descriptive power of such short headings. A detailed table of contents was produced to allow users a more

descriptive view of the handbook. However, the detailed table of contents was long, requiring around 40 pages to view the entire outline. A great deal of effort would be required to develop and maintain the detailed table of contents as the handbook was being created. Therefore, the designers were interested in knowing if users would prefer using the detailed table of contents. A search facility is also planned for the handbook but was not included at the time of our testing.

Navigation is provided within the handbook in several ways. A toolbar (see Figure 2) provides direct access to the table of contents page via an icon. Arrows allow users to move linearly, both forward and backward, in the handbook. Links at the top of each section allow the user to move up in the hierarchy to the chapter level. The designers were interested in discovering whether users would locate and use these navigational aids.

Each chapter contains an outline which users can use to navigate to particular sections within the chapters. A final question was whether users would find these outlines useful. Figure 3 shows the outline for a chapter of the handbook.

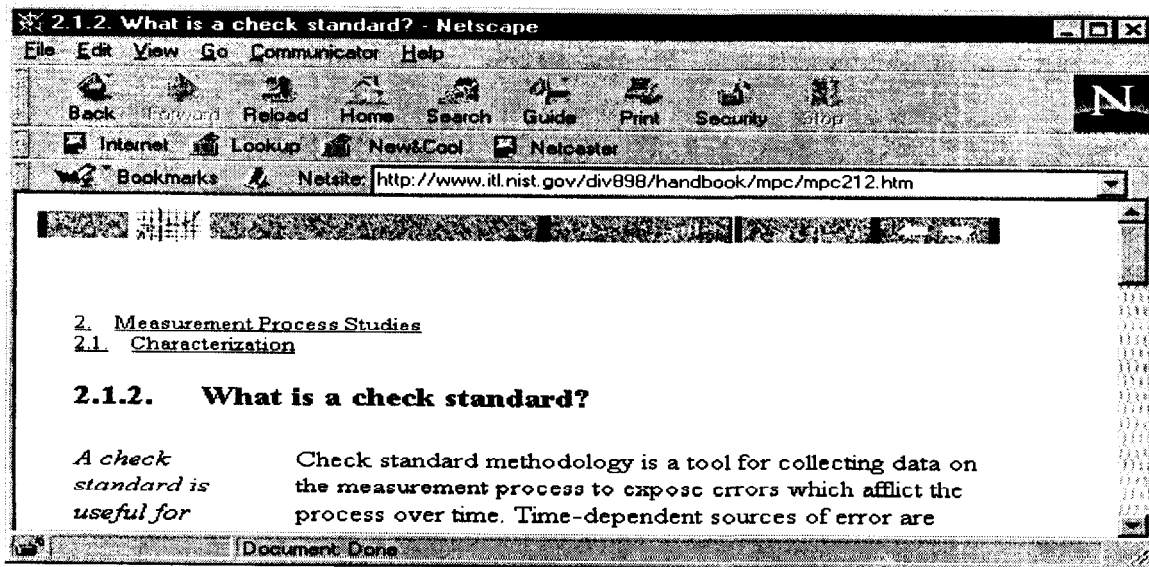


Figure 2. Navigational aids in the handbook

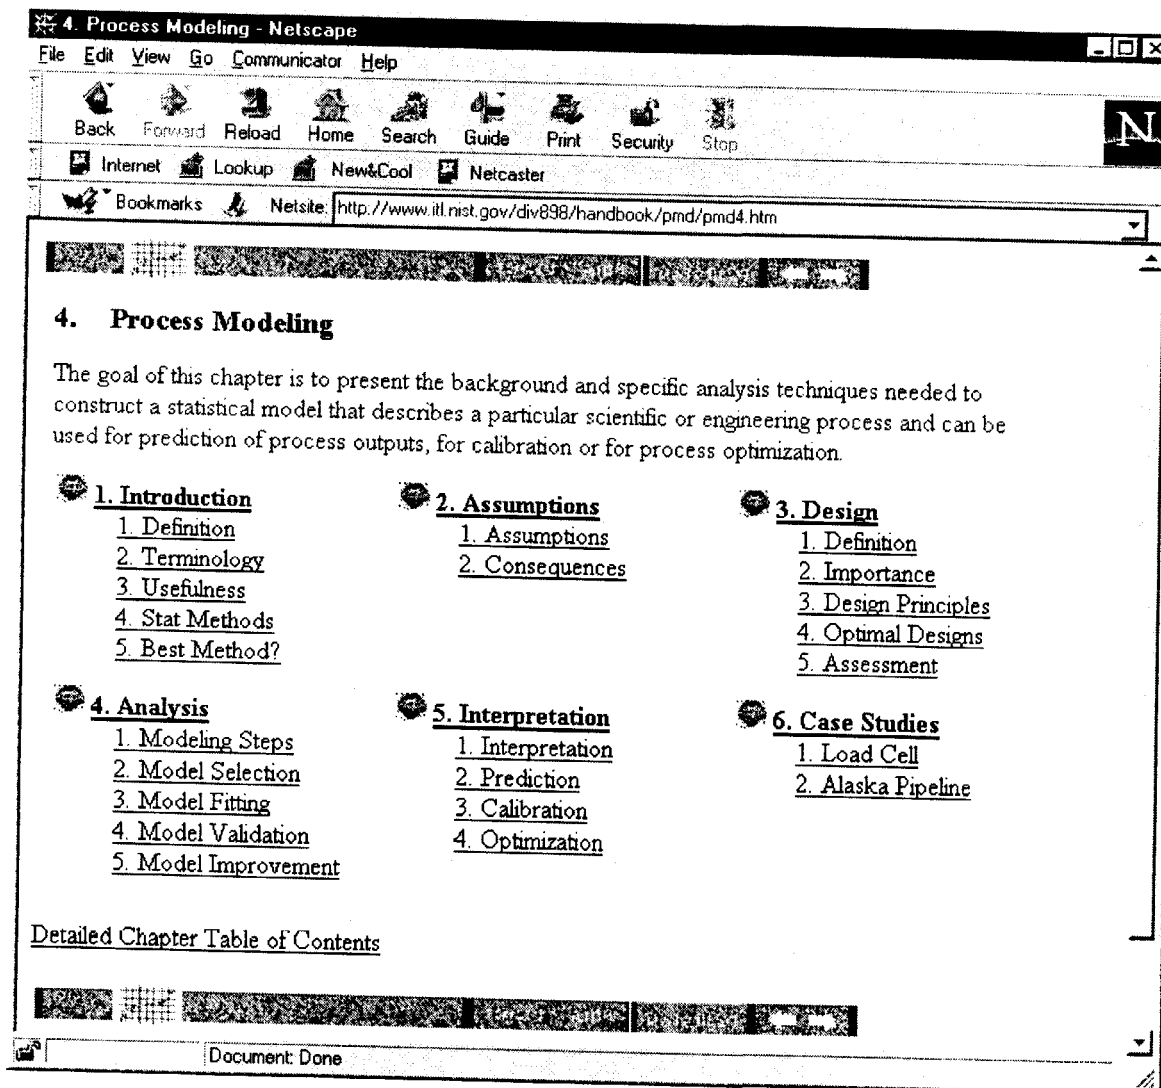


Figure 3. The handbook outline for chapter 4

4.2. Goals of the Usability Research Team

Our goal was to determine the feasibility of using a stand-alone, kiosk-based test to answer the design team questions outlined in the previous section. We needed to meet the requirements we described earlier for a successful kiosk-based test. In designing and conducting this case study, we wanted to identify the potential for tools and techniques we could further develop that would be beneficial to usability professionals for testing on-line books.

We were able to conduct the kiosk usability test twice. First, the kiosk was setup in the poster and demonstration room at an engineering conference held at NIST. During

this conference, the usability research team was present to monitor the test but provided help to users only when absolutely necessary. We also had the opportunity to repeat the experiment the following week during a conference at another location. The usability research team was not present during this experiment. The design team was there and again, monitored the experiment giving help only when absolutely necessary. These two experimental situations allowed us to analyze the performance of the subjects to determine whether there was any significant difference between the two conditions and to determine if we had successfully created a "stand-alone" evaluation.

5. Implementation of the Case Study

The first challenge was to produce a test that was short enough that we could attract participants during breaks at a conference. While kiosk-based or remote tests are not necessarily limited to conferences, users will participate in these types of tests in an "unplanned" fashion. That is, a potential user will walk by a kiosk or discover the test on the web. Therefore, the time required should be short enough that the majority of users will be able to do it immediately, and not have to return later when they have more time. The tasks we wanted users to carryout required users to determine when they had arrived at the correct page. To do this type of testing with the full handbook could take a significant amount of time if

many pages had to be read to determine when the correct one was reached. Another aspect of this study was that the designers wanted feedback early - before all the content of the handbook had been assembled. Therefore, we determined for each task, several alternative paths users could take to reach the desired page. The paths we implemented were the most direct path and the one and two -off paths. That is, users might not select the first link on the direct path but would be able to reach the page by making the correct selection from the page they had first navigated to. We provided immediate feedback on all other paths telling the user that this path was not correct and giving them an easy way of returning to the previous page. While this did not test the readability of the content, it did test the clarity of the links and the navigation tools provided.

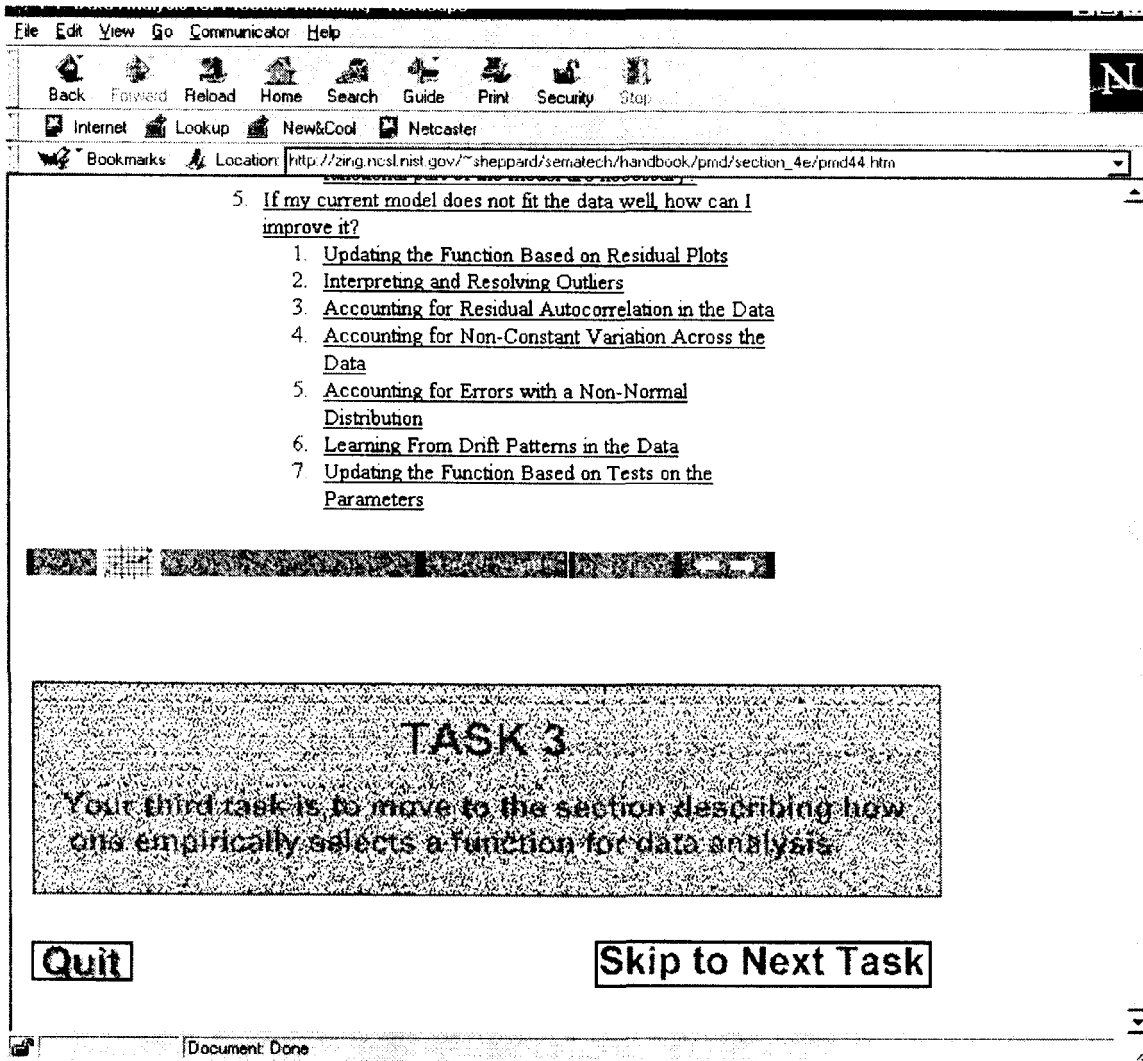


Figure 4. A task displayed on a page of the handbook

The test had to be stand-alone. This meant that we had to make sure that all directions were extremely clear and that we integrated the directions into the experiment materials well but did not confuse users as to what belonged to the handbook and what belonged to the experiment. We also had to provide escape mechanisms for users from all tasks at all times, either skipping to the next task or exiting the experiment completely. This had to be done carefully so that our log files were maintained. We also decided to provide a reminder of the task the user was to do at many points along the path. Otherwise, we were concerned that users would forget the task and have to backtrack to reread it, thereby confounding the navigation required for the task. Figure 4 shows a page from the handbook with the user task displayed on it. We used a very distinctive color to denote the task description and we told users in the initial description of the study about the location of the task description at the bottom of pages. Figure 4 also shows the QUIT and SKIP TASK buttons that the user can use to exit from the usability test or to skip this particular task.

We used WebVIP, one of the tools we had developed in our WebMetrics suite to automate data collection. This tool allowed us to capture each link the user followed along with a time stamp. In addition, our tool captures the use of the back and forward button in the browser should users use that for navigation. Log files were automatically created for each user. Demographic data and rating data that we requested from the users was also written to these log files.

We asked users to do five navigation tasks. The tasks were presented to all users in the same order as the tasks required them to move further down into sections of the handbook. This reflects one way in which the handbook would be used so we were not overly concerned about not randomizing the tasks. We felt that doing so would only serve to confuse the participants. The tasks and the question we were addressing with that task are shown in Figure 5.

Task	Question addressed
Move to the chapter where you would find information about process modeling	Are one word chapter headings usable? Do users use detailed table of contents?
Move to the section on data analysis for process modeling. Move to the section describing how one empirically selects a function.	Can users use the chapter/section outlines to locate sections?
Move to section 4.4.3	Can users use the hierarchical links provided to navigate to other sections and chapters?
Move to section 4.4.3.1 in as few clicks as possible.	Can users use the handbook provided arrows to move sequentially through the handbook?

Figure 5: The tasks used in the case study

6. Results

The results of the experiment are presented in two sections. First we describe what we learned about the usability of the navigation aspects of the on-line book. In the next section we discuss what we learned about conducting kiosk-based tests. For each task we compared the paths that users took to the optimal paths, using the log file data. For this analysis, the comparison was done manually. However,

we are currently developing a visualization tool to display user paths.

6.1. Usability of the handbook navigation

We had fourteen conference attendees participate in the navigation experiment in our first session. In the second session conducted the following week by the design team, there were 11 participants. Figure 6 presents the cumulative results for the tasks. All subjects successfully completed all tasks, with the exception of one subject who skipped a task.

Task	Results
1. Move to the chapter where you would find information about process modeling	68% of the users went directly to the correct chapter heading.
2. Move to the section on data analysis for process modeling.	60% of the users went directly to the correct section. (4.4) Another 20% of the users went to section 4.4.1.
3. Move to the section describing how one empirically selects a function.	52% of the users went directly to the correct section.
4. Move to section 4.4.3	58% of the users went directly to this section. Half used the arrows in the toolbar to navigate to this section and the other half used the hierarchical links at the top of the section.
5. Move to section 4.4.3.1 in as few clicks as possible.	37.5% of the users used the arrows to move to this section.

Figure 6: Results of the navigation experiment

The fact that all users (with the exception of one who quit after a few tasks) were able to complete each task showed that the navigation was successful. However, the results shown in Figure 6 suggest that the optimal methods for navigation could be improved. The one-word chapter headings work reasonably well, with one exception. One chapter was labeled "Explore." Some of our users tried this chapter for the first task. We were concerned that the wording of this chapter implied to users that they could explore the handbook by going to that chapter.

The second task was designed to test if users were able to use the chapter outlines to navigate to sections within the chapter. We had intended for subjects to go to section 4 in chapter 4. Sixty percent of the users did. However, another twenty percent of the users went to section 4.4.1. We do not know whether this was due to the wording of the task or if subjects did not expect to find an actual section at 4.4. We have recommended further laboratory testing using think-aloud protocols (ref) to investigate this.

The third and fourth tasks showed that users were able to use the section and chapter links successfully. Twenty nine percent of the users in task 3 went first to section 4.4.2. We are concerned, however, that users do not fully understand which sections contain direct links to

other sections and which do not. We have recommended using some notation to make this model clear.

The last task was designed to test users' use of the forward arrow on the handbook toolbar. Interestingly, while some users did use the arrows in task four, they did not necessarily use the arrows in task 5. Just over thirty seven percent of the users used the forward arrow immediately to move to the desired section. We have given the design team some recommendations to clarify the functionality of the arrows for users.

While we did not have a task that explicitly tested the use of the detailed table of contents, it was available for users during the tasks. During both sessions, only three users went to the detailed table of contents. Only one of those used the outline to navigate the desired section; the other two immediately went back and tried another path. We have recommended to the designers that the detailed table of contents be made a low priority and that a good internal search engine be developed.

We asked subjects to give ratings about using the navigation aids in the handbook. We used a 7 point scale with 1 being very easy and 7 being very difficult. Figure 7 shows the questions and the average responses we received.

Question	Average
Navigating using the outline was	3.0
Using the arrows provided in the handbook toolbar was	3.2
Using the section navigation links at the top of the page was	3.1
Overall, navigation in the handbook was	3.0

Figure 7: Participant's ratings on navigation

We also collected times for the tasks. The times have no particular value in this particular experiment but can be used to test redesigns to see if subjects can navigate more quickly.

6.2. Lessons learned for kiosk-based testing

Setting up this test was extremely time-consuming. In order to minimize the reading that our participants had to do we copied the handbook site for each task we asked them to do and "operationalized" only the most feasible paths. For example, looking back at Figure 1, subjects were expected to go to the button for Chapter 4. The links for the remaining chapters were replaced with pages that told the subject this was not the correct path. The detailed table of contents was a feasible path so this link remained. However, all links within the detailed table of contents except chapter 4 were replaced with our "dummy" page. However, we had to make as many copies of the dummy page as there were links in order to track the specific links that the user visited.

In several places we inadvertently mislead users by not implementing paths that were actually feasible. In several cases, users started using the forward arrow to reach a section but stopped part way through and tried another path from there. We had not anticipated this and while we had allowed for users to use the arrows, we had not provided the necessary links for them to change tactics in midstream. And, as the tasks took us deeper into the handbook, the number of alternative paths increased making the later tasks much more difficult to setup.

The wording of the tasks is extremely important. There is no opportunity for users to ask questions and results can be skewed depending on their interpretation of the task. We mentioned in the previous section that we were unsure about our results for task 2. We are not sure if users misinterpreted our task directions or if users did not expect to find actual content at section 4.4. We did four pilot studies and many revisions on our own to make sure the wording of our directions was understandable.

It is extremely important to explain exactly, but concisely, the overall directions and the directions for the first task. Once participants started they felt comfortable but it was difficult for several to understand the concept at first. Several participants thought they could just explore the handbook as they wished. They may also have been confused because our kiosk was adjacent to the group demonstrating the handbook.

It was important to our participants to know approximately how long the experiment would take. During our pilot subjects we had recorded the times and we put a time estimate on the beginning screens. We had also implemented a short experiment for the case study design of the handbook [6] that took approximately twice the time as the navigation experiment (1 minutes versus 8 minutes) and some of the subjects choose the navigation portion based strictly on the time.

On a more positive note, we were very encouraged by the number of conference attendees who participated. Counting the number of participants in both the navigation and case study experiments we had 26 subjects in the first session and 15 subjects in the second session we conducted. The first session was conducted during conference breaks during one day only. Moreover, subjects were happy to participate and several told us that it was fun.

The tests were relatively error free. In particular, the navigation portion went smoothly. We did have some problems with the case study portion of the experiment as users sometimes clicked too quickly on steps in the handbook, causing the software to get out of synch. We are recommending the use of kiosk-based testing for navigation types of experiments but we do not advise using this type of experiment for running software.

We were also interested to see if there was any different between the two experimental conditions. In the first condition, members of the usability team were present to oversee the kiosk but we gave help only when absolutely needed (mostly for the software portion of the experiment). In the second condition, the design team

was present. We had given them instructions on when to intervene and when not to and they had also been present to watch the first session. We did several single factor analyses of variance to determine if we had the same user population. We found no significant differences for any of the task times or ratings. We did however, find a significant difference in performance for one of the five tasks: task 2 (involving selection of a section from a chapter outline). In this case the group in the session monitored by the design team performed better than the

group monitored by the usability team. Figure 8 shows the analyses for the performance of the five tasks. At this point, we have no way of knowing the attribution of this performance difference - a quieter setting for the kiosk, more help from the monitors, or just more savvy subjects. Interestingly, the average ratings for overall navigation were lower for the design team monitored group than for the group monitored by the usability team (2.4 versus 3.2).

Task 1	Task 2	Task 3	Task 4	Task 5
Degrees of freedom: 24 F = 1.7 F (critical)= 4.27 P value = .20	Degrees of freedom: 22 F = 4.379 F (critical)= 4.324 P value = .048 Significant	Degrees of freedom: 22 F = 3.75 F (critical)= 4.324 P value = .066	Degrees of freedom: 23 F = 1.34 F (critical)= 4.30 P value = .25	Degrees of freedom: 23 F = 4.01 F (critical)= 4.30 P value = .057

Figure 8: Performance differences analyses for the two sessions in the case study

7. Implications for tools and techniques

We spent considerable effort producing the test materials for the kiosk-based test. First of all, duplicating the actual web site and "operationalizing" only selected paths was extremely time consuming. In a small site, this would not be as difficult nor as time consuming. We are currently investigating ways to help usability professionals automate some of this.

Using WebVIP to automatically collect the time stamped links that users visited saved a great deal of time. We are also investigating ways to speed up the analysis by automatically comparing feasible paths to paths that users actually take. In this case, that analysis was carried out by hand.

The lessons learned exist in the form of guidelines for designing and conducting kiosk-based tests:

- Provide a time estimate for the length of time the users will need to complete the evaluation. This times should definitely be less than 15 minutes and preferably less than 10 minutes.
- Pilot test your instructions until users have no questions and are able to correctly follow them.
- Provide a graceful exist from every step of the evaluation.
- If you need human-subject board approval, make sure a signed consent form is waived. Include a consent agreement that subjects can read and agree to in your instructions.

- Make sure users cannot reach an error state in the kiosk. No matter what actions the user does, the kiosk software must always recover on its own. This includes correctly managing the data collection.
- The kiosk software must correctly reset after each user and be ready for the next user.
- If the evaluation is conducted at a physical kiosk, provide comfortable seating for participants.
- Provide enough context information so users understand the domain and the information you're interested in collecting before they start the experiment. You might consider giving users an example question or task to perform before the actual evaluation starts.

We are interested in doing more case studies on different types of web sites and documents to determine some guidelines for the types of applications and web sites that are suited to kiosk-based testing and the type of data that can be collected via kiosk-based testing.

8. Conclusion

Kiosk-based testing has many advantages. A large number of users can be tested in a relatively short amount of time. If the testing is done in conjunction with conferences or even as a standing public kiosk (placed in the cafeteria or employee lounge area), it serves also as publicity for an effort. Users appreciate attempts to gather feedback from them about the usability of software they will be asked to use.

Kiosk-based testing was very useful in our navigation questions. The information we obtained, paths used for specific tasks and rating information, was useful and answered questions that the design team had posed. We

mentioned that we tried to use this type of testing for an experiment involving the method for running software from the handbook. This was much less successful. We were unable to control many of the conditions and participants would not have been able to recover and proceed with the experiment had we not been present during the case study.

Another discussion point is where in the usability cycle kiosk-based testing fits and how it fits with more traditional usability testing. Kiosk-based testing, of course, can be used to obtain many different types of information. For example, preference data about screen designs and icons can be gathered. Users could be asked to match items to categories to evaluate the developers' structural organization of a web site. As an aside, our WebCAT tools could be used to develop such an evaluation. Performance and satisfaction data for small tasks such as locating information on a small web site or filling in information on a form is easy to obtain. These examples all collect quantitative data, however. If developers are interested in obtaining qualitative data, they should conduct face-to-face studies or remote evaluations using audio and video data collection. This allows developers to use methods such as think-aloud protocols and interviews to determine not only what users did, but why. Kiosk-based evaluation is useful in obtaining more quantitative data to accompany the qualitative data.

The main drawback to kiosk-based testing is the amount of work needed to setup such an evaluation. WebVIP

and WebCAT tools are a good example of tools to ease this burden.

We will be investigating tool design as well as carrying out more kiosk-based case studies in the future.

9. References

1. Hartson, H., Castillo, J. and Kelso, J., Remote Evaluation: The Network as an Extension of the Usability Laboratory, (1996). *Proceedings of CHI 96*, pp. 228-235. New York: ACM Press.
2. The NIST WebMetrics Tool Suite, <http://zing.ncsl.nist.gov/~webmet/>
3. Salomon, G. (1990). Designing Casual-Use Hypertext: The CHI'89 InfoBooth, *Proceedings of CHI '90*, pp. 451-458. New York: ACM Press. pp. 451-458.
4. Scholtz, J., Downey, L. (1998). Methods for Identifying Usability Problems with Web Sites. To appear in EHCI'98 Proceedings.
5. Scholtz, J., Laskowski, S, and Downey, L.. (1998). Developing Usability Tools and Techniques for Designing and Testing Web Sites, 4th Annual Human Factors and the Web Conference, June 5, ATT, Basking Ridge, NJ.
6. Scholtz, J. (1998) Kiosk-Based User Testing of On-Line Books, to appear in Proceedings of SIGDOC conference. Sept, 1998, Quebec City, Canada.
7. Vora, P. (1998). Designing for the Web: A survey, *Interactions*, Vol. 5 (3), pp.13-30.